

Claims

What is claimed is:

1. A method for optimizing processing of a request, the request having elements, the method including:

5        selecting an access path for the request taking into consideration a stored actual request element cost;  
         processing the request using the selected access path, producing one or more actual request element costs; and  
         storing the one or more actual request element costs.

10    2. The method of claim 1 where the request includes one or more predicates, each predicate referencing one or more relations and selecting an access path includes:

         identifying the relations in the request;

         identifying one or more access paths for each relation;

         extracting the predicates from the request;

15        for each predicate,

         associating the predicate with the one or more access plans identified for the one or more relations referenced in the predicate;

         estimating the cost of one or more access paths associated with the predicate; and

         for each access path,

20        selecting the cheaper of the estimated access path cost and an actual access path cost, if one exists.

3. The method of claim 2 where estimating the cost of one or more access paths includes:

         retrieving estimated costs stored in a data dictionary.

4. The method of claim 2 where estimating the cost of one or more access paths includes:

25        performing selectivity costing based on one or more of selectivity, cardinality and statistics.

5. The method of claim 2 where selecting an access path further includes:

         for each predicate,

         selecting the lowest cost access path from among the access paths associated with each predicate.

6. The method of claim 1 where requests are sorted into workload groups and the method further includes:

categorizing the actual request element costs according to the workload group to which the request belongs.

5 7. The method of claim 6 where selecting an access path for the request taking into consideration a stored actual request element cost includes:

taking into consideration the categorized actual request element cost, if it exists, for the workload group to which the request belongs.

8. The method of claim 1 where storing the one or more actual request element costs includes:

10 storing the one or more actual request element costs in a cache.

9. The method of claim 1 where storing the one or more actual request element costs includes:

backing up the stored one or more actual request element costs from the cache to a query capture data base.

10. The method of claim 1 further including:

15 assigning the request to an one of a plurality of workload groups; and  
where selecting an access path for the request includes taking into consideration a stored actual request element cost if the one of the plurality of workload groups is identified for such processing; and

20 where selecting an access path for the request does not include taking into consideration a stored actual request element cost if the one of the plurality of workload groups is not identified for such processing.

11. The method of claim 10 where storing the one or more actual request element costs includes:

performing a workload analysis to associate each actual request element cost with the workload group to which the request is assigned.

25 12. The method of claim 11 where performing the workload analysis includes:

performing at least a portion of the workload analysis off-line.

13. The method of claim 11 where performing the workload analysis includes:

performing at least a portion of the workload analysis in real time.

14. The method of claim 1 further including:

creating one or more workload groups; and  
assigning the request to one of the one or more workload groups.

15. A method for optimizing processing of a workload group, the workload group including one or  
5 more requests, each request having elements, the method including:

selecting an access path for a request taking into consideration a stored actual request element  
cost, categorized by workload group;  
processing the request using the selected access path, producing one or more actual request  
element costs; and  
10 storing the one or more actual request element costs, categorized by workload group.

16. The method of claim 15 where the selected request includes one or more predicates, each predicate  
referencing one or more relations and selecting an access path includes:

identifying the relations in the request;  
identifying one or more access paths for each relation;  
15 extracting the predicates from the request;  
for each predicate,  
associating the predicate with the one or more access plans identified for the one or more  
relations referenced in the predicate;  
estimating the cost of one or more access paths associated with the predicate; and  
20 for each access path,  
selecting the cheaper of the estimated access path cost and an actual access path cost, if one  
exists.

17. The method of claim 16 where estimating the cost of one or more access paths includes:  
retrieving estimated costs stored in a data dictionary.

25 18. The method of claim 16 where estimating the cost of one or more access paths includes:  
performing selectivity costing based on one or more of selectivity, cardinality and statistics.

19. The method of claim 16 where selecting an access path further includes:

for each predicate,

selecting the lowest cost access path from among the access paths associated with each predicate.

5 20. The method of claim 15 where the method further includes:

categorizing the actual request element costs according to the workload group to which the request belongs.

21. The method of claim 20 where selecting an access path for the request taking into consideration a stored actual request element cost includes:

10 taking into consideration the categorized actual request element cost, if it exists, for the workload group to which the request belongs.

22. The method of claim 15 where storing the one or more actual request element costs includes:

storing the one or more actual request element costs in a cache.

23. The method of claim 15 where storing the one or more actual request element costs includes:

15 backing up the stored one or more actual request element costs from the cache to a query capture data base.

24. The method of claim 15 further including:

where selecting an access path for the request includes taking into consideration a stored actual request element cost if the one of the plurality of workload groups is identified for such processing; and

20 where selecting an access path for the request does not include taking into consideration a stored actual request element cost if one of the plurality of workload groups is not identified for such processing.

25. The method of claim 24 where storing the one or more actual request element costs includes:

25 performing a workload analysis to associate each actual request element cost with the workload group to which the request is assigned.

26. The method of claim 25 where performing the workload analysis includes:

performing at least a portion of the workload analysis off-line.

27. The method of claim 25 where performing the workload analysis includes:

performing at least a portion of the workload analysis in real time.

28. A computer program, stored on a tangible storage medium, for use in optimizing processing of a request, the request having elements, the program including executable instructions that cause a computer to:

select an access path for the request taking into consideration a stored actual request element cost;

process the request using the selected access path, producing one or more actual request element costs; and

store the one or more actual request element costs.

29. The computer program of claim 28 where the request includes one or more predicates, each predicate references one or more relations and when selecting an access path the computer:

identifies the relations in the request;

identifies one or more access paths for each relation;

extracts the predicates from the request;

for each predicate,

associates the predicate with the one or more access plans identified for the one or more relations referenced in the predicate;

estimates the cost of one or more access paths associated with the predicate; and

for each access path,

selects the cheaper of the estimated access path cost and an actual access path cost, if one exists.

30. The computer program of claim 29 where when estimating the cost of one or more access paths the computer:

retrieves estimated costs stored in a data dictionary.

31. The computer program of claim 29 where when estimating the cost of one or more access paths the computer:

performs selectivity costing based on one or more of selectivity, cardinality and statistics.

32. The computer program of claim 29 where when selecting an access path the computer further:  
for each predicate,

selects the lowest cost access path from among the access paths associated with each  
predicate.

5 33. The computer program of claim 28 where requests are sorted into workload groups and the  
computer program further includes executable instructions that cause the computer to:

categorize the actual request element costs according to the workload group to which the  
request belongs.

34. The computer program of claim 33 where when selecting an access path for the request taking into  
10 consideration a stored actual request element cost the computer:

takes into consideration the categorized actual request element cost, if it exists, for the  
workload group to which the request belongs.

35. The computer program of claim 28 where when storing the one or more actual request element  
costs the computer:

15 stores the one or more actual request element costs in a cache.

36. The computer program of claim 28 where when storing the one or more actual request element  
costs the computer:

backs up the stored one or more actual request element costs from the cache to a query capture  
data base.

20 37. The computer program of claim 28 further including executable instructions that cause a computer  
to:

assign the request to an one of a plurality of workload groups;

where when selecting an access path for the request, the computer takes into consideration a  
stored actual request element cost if the one of the plurality of workload groups is  
25 identified for such processing; and

where when selecting an access path for the request, the computer does not include take into  
consideration a stored actual request element cost if the one of the plurality of workload  
groups is not identified for such processing.

38. The computer program of claim 37 where when storing the one or more actual request element costs the computer:

performs a workload analysis to associate each actual request element cost with the workload group to which the request is assigned.

5 39. The computer program of claim 38 where when performing the workload analysis the computer:  
performs at least a portion of the workload analysis off-line.

40. The computer program of claim 38 where when performing the workload analysis the computer:  
performs at least a portion of the workload analysis in real time.

41. The computer program of claim 28 further including executable instructions that cause a computer  
10 to:  
create one or more workload groups; and  
assign the request to one of the one or more workload groups.

42. A database system including:

a massively parallel processing system including:

15 one or more nodes;  
a plurality of CPUs, each of the one or more nodes providing access to one or more  
CPUs;  
a plurality of data storage facilities each of the one or more CPUs providing access to  
one or more data storage facilities;  
20 a process for execution on the massively parallel processing system for optimizing processing  
of a request, the request having elements, the process including:  
selecting an access path for the request taking into consideration a stored actual request  
element cost;  
processing the request using the selected access path, producing one or more actual request  
25 element costs; and  
storing the one or more actual request element costs.

43. The database system of claim 42 where the request includes one or more predicates, each predicate referencing one or more relations and selecting an access path includes:

identifying the relations in the request;

identifying one or more access paths for each relation;

5 extracting the predicates from the request;

for each predicate,

associating the predicate with the one or more access plans identified for the one or more relations referenced in the predicate;

estimating the cost of one or more access paths associated with the predicate; and

10 for each access path,

selecting the cheaper of the estimated access path cost and an actual access path cost, if one exists.

44. The database system of claim 43 where estimating the cost of one or more access paths includes: retrieving estimated costs stored in a data dictionary.

15 45. The database system of claim 43 where estimating the cost of one or more access paths includes: performing selectivity costing based on one or more of selectivity, cardinality and statistics.

46. The database system of claim 43 where selecting an access path further includes:

for each predicate,

selecting the lowest cost access path from among the access paths associated with each predicate.

47. The database system of claim 42 where requests are sorted into workload groups and the process further includes:

categorizing the actual request element costs according to the workload group to which the request belongs.

25 48. The database system of claim 47 where selecting an access path for the request taking into consideration a stored actual request element cost includes:

taking into consideration the categorized actual request element cost, if it exists, for the workload group to which the request belongs.



49. The database system of claim 42 where storing the one or more actual request element costs includes:

storing the one or more actual request element costs in a cache.

50. The database system of claim 42 where storing the one or more actual request element costs includes:

backing up the stored one or more actual request element costs from the cache to a query capture data base.

51. The database system of claim 42 where the process further includes:

assigning the request to an one of a plurality of workload groups; and

where selecting an access path for the request includes taking into consideration a stored actual request element cost if the one of the plurality of workload groups is identified for such processing; and

where selecting an access path for the request does not include taking into consideration a stored actual request element cost if the one of the plurality of workload groups is not identified for such processing.

52. The database system of claim 51 where storing the one or more actual request element costs includes:

performing a workload analysis to associate each actual request element costs with the workload group to which the request is assigned.

53. The database system of claim 52 where performing the workload analysis includes:

performing at least a portion of the workload analysis off-line.

54. The database system of claim 52 where performing the workload analysis includes:

performing at least a portion of the workload analysis in real time.

55. The database system of claim 42 where the process further includes:

creating one or more workload groups; and

assigning the request to one of the one or more workload groups.

56. The database system of claim 42 further comprising a plurality of additional process for optimizing processing of a request.